

REMARKS

The above amendments are submitted within the three-month period for response to the Final Office Action mailed May 16, 2007 and in connection with a Request for Continued Examination (RCE). Authorization for a Credit Card charge of the \$790.00 requisite fee is hereby included in the Electronic Fee Sheet attached. Reconsideration and allowance of all pending claims are respectfully requested.

In the subject Office Action, claim 7 was objected-to by the Examiner. Furthermore, claim 1 was rejected under 35 U.S.C. § 102(b) as being anticipated by U.S. Patent No. 6,556,990 to Lane. In addition, claims 2, 17 and 19 were rejected under 35 U.S.C. § 103(a) as being unpatentable over Lane in view of U.S. Patent No. 6,785,677 to Fritchman, and claim 18 was rejected under 35 U.S.C. § 103(a) as being unpatentable over Lane in view of Fritchman and further in view of U.S. Patent No. 5,845,276 to Emerson et al.

Applicants respectfully traverse the Examiner's rejections to the extent that they are maintained. Applicants have amended claims 1, 17 and 19 herein, and Applicants respectfully submit that no new matter is being added by the above amendments, as the amendments are fully supported in the specification, drawings and claims as originally filed. Applicants also note that the amendments made herein are being made only for facilitating expeditious prosecution of the aforementioned claimed subject matter. Applicants are not conceding in this application that the originally claimed subject matter is not patentable over the art cited by the Examiner, and Applicants respectfully reserve the right to pursue this and other subject matter in one or more continuation and/or divisional patent applications.

As an initial matter, Applicants wish to thank the Examiner for the consideration extended in the telephonic interview conducted between the Examiner, the Examiner's supervisor and Applicants' representative on August 9, 2007. In the interview, proposed amendments to the claims were discussed, and the Examiner indicated that the proposed amendments (which are identical to the amendments made herein, overcame the rejections.

The Examiner did request that Applicants provide citations from the specification supporting the new claim language.

Now turning to the subject Office Action, and initially to the Examiner's objection to claim 7, the Examiner will note that Applicants have corrected the status identifier of the claim from "Previously Presented" to "Withdrawn", as required by the Examiner. Withdrawal of the objection is therefore respectfully requested. Applicants do note, however, that upon allowance of claim 1, claims 3-15 should be reinstated and reconsidered as being dependent upon an allowed claim.

Now turning to the art-based rejections, and specifically to the rejection of independent claim 1, this claim as amended recites a method for indexing a database table that comprises a column of values. The claim, as amended, recites identifying a plurality of substrings, each substring comprising one or more characters; and for each substring, creating a corresponding index over the column such that a plurality of substring-specific indices, each corresponding to a particular substring from among the plurality of substrings, are created for the same column. Each corresponding index includes a respective entry for each value in the column, and the respective entry for a particular value is indicative of whether the corresponding substring occurs in the particular value.

Claim 1 has specifically been amended to clarify that a plurality of substring-specific indices, each corresponding to a particular substring from among the plurality of substrings, are created for the same column. As such, Applicants have clarified that the claimed indices are (1) substring-specific; and (2) created for the same column. Support for this amendment may be found, for example, at page 8, lines 2-4:

As mentioned above, the embodiments discussed hereinafter utilize a database engine and optimizer framework that support creation of *a series of profiling indices over a column* having character string data, such as a traditional "varchar" data type (*emphasis added*).

Additional support may be found in Figs. 4A-4B, and the supporting disclosure at page 12, line 15 to page 13, line 20. Fig. 4A, for example, illustrates a series of profiling indices 405, 407, 409 that are substring-specific, as each index 405, 407, 409 corresponds to a single substring 404, 406, 408. Table 400 in fact is shown with 256 different

substring-specific indices corresponding to the 256 ASCII characters (*see* p. 12, line 23-24, “table 400 depicts 256 different sub-strings, each consisting of a single character and each having its own profiling index.” (*emphasis added*)). Consequently, the specification adequately supports the concept that the claimed indices are “substring-specific.”

It is also evident that the claimed indices are all created on the same column. As discussed at page 13, lines 3-10, for example, all of indices 405, 407, 409 are built over the same column 402 (p. 13, lines 7-8, “in the table 400, each profiling index 405, 407, 409 is built over the Field2 (i.e., column 402).”)

As discussed in the interview, the fact that the claimed indices are substring-specific and all created for the same column often enables the indices to substantially reduce the number of records that need to be retrieved and searched when executing a query, particularly when executing a query that uses a LIKE predicate. In the example shown in Fig. 4A, for example, it can be readily ascertained that record 1 includes the letter “b” in column 402, but not either of letters “a” or “Z.” Likewise, record 2 includes letters “a” and “b”, but not “Z,” and record 10,000 includes all three letters. If, therefore, a query is processed where the letter “Z” is included in the query, records 1 and 2 can be eliminated as possible matches simply by accessing the “Z”-specific profiling index 409, and without having to retrieve the records themselves.

In rejecting claim 1, the Examiner relies on Lane. As discussed previously, however, Lane does not disclose or suggest a plurality of substring-specific indices created over the same column. Instead, Lane discloses two separate embodiments of data structures that may be used to locate substrings in a table.

The first embodiment, shown in Fig. 3 and described at col. 4, line 65 to col. 5, line 10, relies on a single “sub-string index” labeled 305. The sub-string index has a first column with identifiers of rows in the table, and a second column with sub-strings associated with the strings in the rows in the table. Thus, as can be seen in Fig. 3, index 305 includes three entries mapping Row 001 of the table (which contains the string “ABCDE”) to three sub-strings “ABC”, “BCD”, and “CDE”.

Consequently, this embodiment discloses a single index. Furthermore, this single index includes entries that correspond to multiple, different substrings. Claim 1, on the other hand, recites a plurality of substring-specific indices, with each index corresponding to a particular substring. This embodiment of Lane therefore does not disclose either substring-specific indices (as multiple substrings are represented in a single index) nor multiple indices created over the same column (as only a single index is disclosed).

The second embodiment of Lane, also shown in Fig. 3, and described at col. 5, lines 11-33, relies upon a "sub-string table" labeled 308, which includes a first column with strings from the target column of the table and a second column with sub-strings associated with the strings in the first column. Thus, as may be seen in Fig. 3, three rows are provided in the sub-string table for the string "ABCDE", with each row mapping the string to a different sub-string, here "ABC", "BCD", and "CDE".

Two points are worth mentioning with respect to sub-string table 308. First, the sub-string table is identified by Lane as "table" and not an "index." Thus, the sub-string table 308 cannot be interpreted as one of a plurality of indices within the context of claim 1. Second, sub-string index 305 and sub-string table 308 are described as being from alternative embodiments, i.e., they are never used together in the same embodiment, as they are alternative ways of performing the same function.

In addition, even if table 308 were analogized to an index, it would still fail to anticipate the indices recited in claim 1. As with index 305, table 308 discloses a single data structure. Furthermore, this single data structure includes entries that correspond to multiple, different substrings. Claim 1, on the other hand, recites a plurality of substring-specific indices, with each index corresponding to a particular substring.

The second embodiment illustrated in Fig. 3 also includes two indices 306 and 310. Index 306 is created on the target column of the primary table, as is used to map strings in the target column to row identifiers in the table. Thus, for string "ABCDE", index 306 includes an entry identifying row 001 as containing that string.

Index 310, on the other hand, is created on the sub-string column of sub-string table 308 to map each sub-string in sub-string table 308 to rows in the sub-string table 308. Consequently, for sub-string "CDE", index 310 includes an entry mapping that sub-string to row 002 of sub-string table 308. Thus, to identify rows in the primary table that contain the substring "CDE", Lane first accesses index 310 to locate an entry matching the substring, and when such an entry is found, uses the row number identified in the entry to access the corresponding row in the sub-string table 308. The identified row of sub-string table 308 is then used to locate the complete string "ABCDE" containing the substring, and index 306 is accessed to locate the row number in the primary table that contains the complete string.

Of note, however, as with both index 305 of the first embodiment, and table 308 of the second embodiment, each of indices 306, 310 fails to meet the language recited in claim 1. Each index 306, 310 is a single index. Furthermore, index 306 does not even contain substrings, but rather contains complete strings from the primary table. Index 310, while containing substrings, contains multiple substrings, and thus does not "correspond" to a particular substring, nor is it "substring-specific" as required by claim 1.

In addition, even were all of the indices and tables in the second embodiment collectively considered to represent a set of indices, no individual indices correspond to specific substrings. Consequently, the second embodiment does not disclose a plurality of substring-specific indices created over the same column, as required by claim 1.

It is also important to note that claim 1 recites the specific structure incorporated into each substring-specific index. In particular, claim 1 recites that "each corresponding index includes a respective entry for each value in the column, the respective entry for a particular value being indicative of whether the corresponding substring occurs in the particular value." As shown in Fig. 4A, for example, the "a"-specific index 405, which is specific to the substring "a" 404, includes entries for each value in column 402 that are indicative of whether the corresponding substring occurs in that value. Consequently, entries corresponding to records 1, 2 and 10,000 in index 405 have values of "0," "1," and

“1,” indicating that value 1 in column 402 does not include the substring “a,” while values 2 and 3 in column 402 do include substring “a.”

None of the data structures in Fig. 3 of Lane correspond to the claimed structure of a substring-specific index as recited in claim 1. Index 305 includes in its first column the record number of each record in column 304 of table 302 that includes a particular substring identified in the second column; however, the index does not include an entry for “each value in the column,” as required by claim 1. Table 308 likewise matches substrings with values in a column, but does not necessarily contain an entry for “each value in the column,” as required by claim 1. Index 306 only maps record numbers in table 308 to values in column 304, and index 310 maps substrings to record numbers in table 308, so neither includes an entry for “each value in the column,” as required by claim 1. Consequently, no data structure in Lane corresponds to the claimed structure of the substring-specific indices recited in claim 1.

Applicants therefore respectfully submit that claim 1 is novel over Lane, and the rejection should be withdrawn. Applicants also respectfully submit that claim 1 is non-obvious over Lane, as no objective reason has been proffered establishing that one of ordinary skill in the art would be motivated to modify Lane to incorporate the substring-specific indices recited in claim 1. Certainly Lane, which discloses two different sets of data structures for locating substrings in a table, neither of which meet the language of claim 1, does not provide any such motivation. Applicants therefore respectfully submit that independent claim 1 is patentable over Lane and the other prior art of record. Reconsideration and allowance of claim 1, and of claim 2 which depends therefrom, are respectfully requested. Furthermore, claims 3-15, which depend from claim 1, are also patentable based upon their dependency on claim 1.

Next, with respect to independent claims 17 and 19, each of these claims has been amended in a similar manner to claim 1. Claim 17 as amended now recites “a plurality of substring-specific indices stored within said memory and corresponding to the same column in said table, wherein each index corresponds to a particular one of a plurality of substrings.” Likewise, claim 19 now recites that “a plurality of substring-specific indices,

each corresponding to a particular substring from among the plurality of substrings, are created for the same column.”

As discussed above in connection with claim 1, Lane does not disclose or suggest the claimed substring-specific indices. Furthermore, the secondary reference cited by the Examiner, Fritchman, is relied upon merely for allegedly disclosing "a number of occurrences." Applicants therefore respectfully submit that independent claims 17 and 19 are novel and non-obvious over Lane and Fritchman and the other art of record. Reconsideration and allowance of claims 17 and 19, and of claim 18 which depends therefrom, are therefore respectfully requested.

As a final matter, Applicants traverse the Examiner's rejections of the dependent claims based upon their dependency on the aforementioned independent claims. Nonetheless, Applicants do note that a number of these claims recite other features that further distinguish these claims from the references cited by the Examiner. However, in the interest of prosecutorial economy, these claims will not be addressed separately herein.

In summary, Applicants respectfully submit that all pending claims are novel and non-obvious over the prior art of record. Reconsideration and allowance of all pending claims are therefore respectfully requested. If the Examiner has any questions regarding the foregoing, or which might otherwise further this case onto allowance, the Examiner may contact the undersigned at (513) 241-2324. Moreover, if any other charges or credits are necessary to complete this communication, please apply them to Deposit Account 23 3000.

Respectfully submitted,

August 14, 2007
Date

/Scott A. Stinebruner/
Scott A. Stinebruner
Reg. No. 38,323
WOOD, HERRON & EVANS, L.L.P.
2700 Carew Tower
441 Vine Street
Cincinnati, Ohio 45202
Telephone: (513) 241-2324
Facsimile: (513) 241-6234